

Reinforcement learning from comparisons: Three alternatives is enough, two is not.

Benoît LASLIER

Institut Camille Jordan, Université Claude Bernard Lyon 1,
43 boulevard du 11 novembre 1918, 69622 Villeurbanne, France

Jean-François LASLIER

CNRS and École polytechnique,
91128 Palaiseau, France
jean-francois.laslier@polytechnique.edu

January 25, 2013

Abstract

The paper deals with the problem of finding the best alternatives on the basis of pairwise comparisons when these comparisons need not be transitive. In this setting, we study a reinforcement urn model. We prove convergence to the optimal solution when reinforcement of a winning alternative occurs each time after considering three random alternatives. The simpler process, which reinforces the winner of a random pair does not always converges: it may cycle.

1 Introduction

In a pairwise comparison problem, one is given a set of alternatives, with data about how they compare the ones to the others. In its purest form, on which we focus in the present paper, we simply have, for any pair of distinct alternatives, the information of which one “beats” the other. Such a data set is called a tournament. Basic results on this structure can be found in Moon [18].

For logical as well as practical reasons, binary relations are at the basis of choice theory. Two classical examples are sport competition and majoritarian politics. Many sports involve by definition two players (or teams), so that competition among any number of players must take the form of

the aggregation of pairwise comparisons. In majority voting, a candidate is socially preferred to another if and only if a majority of the voters prefer the former to the latter. More generally, the prevalence of that kind of binary relations can be traced back to specific features of efficient natural languages (Rubinstein [25]).

If a chess player beats all the other players, he or she is clearly the best. If a candidate cannot be defeated under majority rule by any challenger, that “Condorcet” candidate can claim to be the best according to majority rule.¹ But if no alternative beats all the others, it is not clear how to define the best alternatives. The problem of choosing from pairwise comparisons has thus attracted the attention of scholars in various fields, most often from the axiomatic, normative, point of view (David [5]; Fishburn [6]; Laslier [13]; Brandt et al. [1]).

In the present paper we tackle the same problem from an evolutionary perspective instead of an axiomatic one. We consider dynamic processes according to which, at each period in time, a small number (2 or 3) of alternatives are sampled, the tournament is played among these few alternatives, and the winning alternative is reinforced in the sense that it will be sampled with higher probability in the future. Where does such a mechanical adaptive process go? Using a standard urn model, where reinforcing an alternative is adding a colored ball to an urn, we obtain two results.

(i) If one samples three alternatives (distinct or not) at each date, the process is able to discover the optimal solution of the tournament, that is the unique probability distribution p^* which is, in expectation, defeated by no alternative. With probability one, the composition of the sampling urn, which defines the probability p_τ of choosing the various alternatives at time τ , tends to p^* when τ tends to infinity.

(ii) If one samples only two alternatives at each date, the process is not able to discover the optimal solution, unless the solution is degenerated, with one alternative defeating all the others. With probability one, the composition of the sampling urn, which gives the probability p_τ of choosing the alternatives, concentrates on the support of the optimal solution p^* . This means that all the alternatives which are played with zero probability in the optimal solution are chosen with a probability going to 0. However the composition of the urn may cycle, staying away from the optimal solution. In some cases we even prove that it cycles with probability one.

The negative result (ii) echoes known results about the evolutionary instability of mixed equilibria in evolutionary game theory. For instance

¹This observation does not imply that the majority principle is good for Politics.

cycling with probability one is proven by Posch [23] in a reinforcement urn model for 2×2 games.

The positive result (i) seems more original. In a study of imitation processes in Matching Pennies games, Hofbauer and Schlag [11] observe that players end up closer to the equilibrium if they sample several individual before imitating: there is still cycling, but closer to the equilibrium. Our results might be interpreted in the same spirit: learning slower leads to more stability.

The techniques we use to derive these results are standard in the field of adaptive processes with reinforcement; see Pemantle 2007 [22], and they belong to the family of martingales techniques. The main ingredient of the proof is the definition of a well chosen function of the process whose values form a martingale (see (11)). We use the convergence theorem for positive martingales to obtain some global asymptotic information about the process. In the case of three alternatives we get fairly directly the convergence of the process while for the case of two alternatives the convergence theorem has to be complemented with a variance analysis to prove the non-convergence.

The paper is organized as follows. Section 2 introduces the necessary notions about tournaments: definition and notation (2.1), the Markov chain induced by the play of small-size tournaments at each date (2.2), the tournament game which allows to define and to prove existence of the optimal solution (2.3), and some further preliminary material (2.4, 2.5). Section 3 starts by the definition of urns and of the adaptive processes (3.1). Then, in order to illustrate the argument in a simple way, a toy example is introduced and treated according to its deterministic approximation (3.2). The statement and proof of our main result on three-alternatives reinforcement is found in (3.3) and two-alternatives reinforcement is treated in (3.4), before a short conclusion (3.5).

2 Framework

2.1 Tournaments

Let X be a finite set. A *tournament* T on X is a complete and antisymmetric binary relation. For any x and y in X , one and only one of the three possibilities occurs: $x = y$, xTy , or yTx . When xTy we often say that x beats y . Define the sets :

$$T^+(x) = \{y \in X : xTy\}, \quad T^-(x) = \{y \in X : yTx\}. \quad (1)$$

The binary relation T is fixed throughout this paper. It is sometimes easier to use the notation:

$$\max\{x, y\} = \begin{cases} x & \text{if } x T y \text{ or } x = y \\ y & \text{if } y T x. \end{cases}$$

An alternative which beats all other alternatives is called a *Condorcet winner*. A tournament can have a Condorcet winner or not, but cannot have two. The *Top-Cycle* of the tournament T is the smallest (by inclusion) set $Y \subseteq X$ such that:

$$\forall x \in X \setminus Y, \exists y \in Y : y T x.$$

It is easily seen that such a set is unique and reduces to a singleton $\{c\}$ if and only if c is a Condorcet winner. The literature on tournaments and formal political science has shown that the Top-Cycle is usually a very large set (McKelvey [17]), and has proposed many refinements of this set (see [13] for a survey).

2.2 A Markov chain

Let $\Delta(X)$ be the set of probability distributions on X and let $p \in \Delta(X)$. The support of p is denoted by $\text{Supp}(p)$. Given p , define a sequence $(p^{[t]})_{t \in \mathbb{N}}$ of probability distributions on X derived from p in the following way :

$$p^{[0]} = p, \tag{2}$$

$$p^{[t+1]}(x) = p^{[t]}(x) \cdot p(T^+(x) \cup \{x\}) + p^{[t]}(T^+(x)) \cdot p(x), \tag{3}$$

for any $t \geq 0$, for any $x \in X$

The interpretation is that $p^{[t]}$ is the distribution of a random variable $\xi(t) \in X$ such that $\xi(0)$ is chosen at random according to p and then, given that $\xi(t) = x$, $\xi(t+1)$ is the winner (according to T) of the comparison between x and some alternative y randomly chosen in X according to p . Therefore $\xi(t+1) = x$ either because $\xi(t)$ was already equal to x and y was chosen in $T^+(x) \cup \{x\}$ (first term in the above formula), or because $\xi(t)$ was in $T^+(x)$ and x was chosen according to p (second term). We call p the “sampling” probability.

This process is usually considered with p uniform on X (Daniels [4], Ushakov [28], Levchenkov [16], Slutsky and Volij [27], Chebotarev and Shamis [2, 3]). We need the general version because, later in this paper, p will be endogenous.

Given p , the stationary distribution for this finite Markov chain exists and is unique;² we denote it by $p^{[\infty]}$. It is characterized by the fact that $\text{Supp}(p^{[\infty]}) \subseteq \text{Supp}(p)$ and, for any x in $\text{Supp}(p)$,

$$p^{[\infty]}(T^+(x)) \cdot p(x) = p^{[\infty]}(x) \cdot p(T^-(x)). \quad (4)$$

Notice that the inclusion $\text{Supp}(p^{[\infty]}) \subseteq \text{Supp}(p)$ may be strict; indeed, $p^{[\infty]}(x) = 0$ when $p(T^-(x)) = 0$, that is when x beats no alternative in the support of p . More exactly, $\text{Supp}(p^{[\infty]})$ is the Top-Cycle of the restriction of T to $\text{Supp}(p)$; thus $\text{Supp}(p^{[\infty]})$ does not exactly really depends on p but only on $\text{Supp}(p)$. If p has full support, for instance in the usual case where p is uniform, $\text{Supp}(p^{[\infty]}) = TC(T)$.

2.3 The tournament game

The tournament game is the two-player, symmetric, zero-sum game defined by the strategy set X and the payoff function $g(x, y) = +1$ if $x T y$, $g(x, y) = 0$ if $x = y$, and $g(x, y) = -1$ if $y T x$. For $p, q \in \Delta(X)$ two probability distributions on X , write:

$$g(p, q) = \sum_{y \in X} g(x, y) p(x) q(y). \quad (5)$$

From the definition, g is clearly antisymmetric: $g(q, p) = -g(p, q)$.

The tournament game has been studied by graph theorists (Fischer and Ryan [8, 9, 10]) and has more recently attracted attention of computer scientists (Rivest and Chen [24]). As a model of majority voting and two-party electoral competition, it studied in Social Choice theory and formal Political Science (Moulin [19], Myerson [20, 21], Laslier [14, 15]). Remarkably, such a game has a unique equilibrium. Here is the precise result that will be needed in the sequel. (Fisher and Ryan [9] prove this result using linear algebra and Laffond et al. [12] have a direct proof using a parity argument.)

Proposition 1 *There exists a unique $p^* \in \Delta(X)$ such that $g(p^*, q) \geq 0$ for all $q \in \Delta(X)$. This p^* , called the optimal strategy, is also characterized by the following : for all $x \in X$,*

$$\begin{aligned} p^*(x) > 0 &\iff g(x, p^*(x)) = 0 \\ p^*(x) = 0 &\iff g(x, p^*(x)) < 0. \end{aligned}$$

²We state the results in this section without proofs. They are easily derived from elementary theory of finite Markov chains and have already been noticed for p uniform in the mentioned references.

The support of the optimal strategy is called the *Bipartisan Set* of the tournament: $\text{Supp}(p^*) = BP(T)$. This set is a subset of the Top Cycle and the inclusion is often strict. For instance, in totally random tournaments, the Top Cycle contains all the alternatives and the Bipartisan Set contains only half of them (Fisher and Reeves [7]).

2.4 Two formulas

Before we go further and explain the relation between the game optimal strategy and stationary probabilities, it is useful to state two technical formulas. The following lemma describes the probabilities $p^{[2]}$ and $p^{[3]}$, obtained after sampling two or three alternatives with the Markov chain defined in Section 2.2, in term of the payoff function g .

Lemma 2 *For any $x \in X$:*

$$\begin{aligned} p^{[2]}(x) &= p(x) \cdot (1 + g(x, p)), \\ p^{[3]}(x) &= p(x) \cdot \left(1 + \frac{3}{2}g(x, p) + \frac{1}{2}g(x, p)^2 + \sum_{y \in X} p(y)g(x, y)g(y, p) \right). \end{aligned}$$

Proof. First let us notice a useful equality. By definition (1) and (5)

$$g(x, p) = p(T^+(x)) - p(T^-(x)), \quad (6)$$

and, since $p(T^+(x)) + p(T^-(x)) + p(x) = 1$, we get:

$$1 + g(x, p) = 2p(T^+(x)) + p(x). \quad (7)$$

Let a and b be chosen according to p and let $x = \max\{a, b\}$, then:

$$\begin{aligned} p^{[2]}(x) &= \Pr[a = x] \cdot \Pr[b \in T^+(x) \cup \{x\}] + \Pr[a \in T^+(x)] \cdot \Pr[b = x] \\ &= p(x) \cdot (2p(T^+(x)) + p(x)) \\ &= p(x) \cdot (1 + g(x, p)). \end{aligned}$$

For the second formula, let a , b and c be chosen according to p . An alternative x appears as $x = \max\{\max\{a, b\}, c\}$ in the two exclusive cases:

$$\begin{aligned} x &= \max\{a, b\}, \text{ and } x = \max\{x, c\} . \\ x &T \max\{a, b\}, \text{ and } x = c . \end{aligned}$$

In the first line, the event $x = \max\{a, b\}$ has probability $p^{[2]}(x) = p(x) \cdot (1 + g(x, p))$ so the probability of the first case is $p(x) \cdot (1 + g(x, p)) \cdot p(T^+(x) \cup \{x\})$. In the second line, the event $x T \max\{a, b\}$ has probability

$$p^{[2]}(T^+(x)) = \sum_{y \in T^+(x)} p^{[2]}(y) = \sum_{y \in T^+(x)} p(y)(1 + g(y, p)),$$

therefore the probability $p^{[3]}$ is:

$$\begin{aligned} p^{[3]}(x) &= p(x) \cdot (1 + g(x, p)) \cdot \left(p(x) + \sum_{y \in T^+(x)} p(y) \right) + p(x) \sum_{y \in T^+(x)} p(y)(1 + g(y, p)) \\ &= p(x)^2 [1 + g(x, p)] + p(x) \sum_{y \in T^+(x)} p(y) [(2 + g(x, p)) + g(y, p)]. \end{aligned}$$

Using the fact that $\frac{1+g(x,y)}{2}$ is 1 if $y \in T^+(x)$, is 1/2 if $y = x$, and is 0 if not, one finds:

$$\begin{aligned} &\frac{p^{[3]}(x)}{p(x)} \\ &= \sum_y p(y) [(2 + g(x, p)) + g(y, p)] \frac{1 + g(x, y)}{2} \\ &= \frac{1}{2} \sum_y p(y) [(2 + g(x, p)) + g(y, p) + 2g(x, y) + g(x, p)g(x, y) + g(y, p)g(x, y)] \\ &= 1 + \frac{1}{2}g(x, p) + \frac{1}{2}g(p, p) + g(x, p) + \frac{1}{2}g(x, p)^2 + \frac{1}{2} \sum_y p(y)g(y, p)g(x, y) \\ &= 1 + \frac{3}{2}g(x, p) + \frac{1}{2}g(x, p)^2 + \frac{1}{2} \sum_y p(y)g(y, p)g(x, y), \end{aligned}$$

which is the announced formula.

Q.E.D.

2.5 Relation between optimal strategies and stationary probabilities

We first observe that the game optimal strategy p^* satisfies a nice fixed-point property if we take $p^{[1]} = p^*$ as the sampling probability to build the Markov chain, and that only an optimal strategy can be such a fixed point.

Proposition 3 *Let p^* be the optimal strategy for the tournament game, then $(p^*)^{[2]} = (p^*)^{[\infty]} = p^*$. Conversely, let p be such that $p^{[2]} = p$, then p is the optimal strategy for the tournament game restricted to the support of p .*

Proof. By lemma 2, $p^{*[2]}(x) = p^*(x)(1 + g(x, p^*))$, and, by proposition 1, either $p^*(x) = 0$ or $g(x, p^*) = 0$.

Conversely if $p^{*[2]}(x) = p(x) = p(x)(1 + g(x, p))$ then $g(x, p) = 0$ as soon as $p(x) > 0$ and p is the optimal strategy on its support. **Q.E.D.**

3 Learning

With the previous background material in mind, we turn to the main result of this paper. Instead of considering re-sampling at each date according to a constant probability distribution, as is done in the previously described Markov chains, we describe learning processes where winning alternatives are reinforced at the level of the sampling probability. These processes can be implemented with random urns.

3.1 Choice by reinforcement

An *urn* on X is a list n of strictly positive integers $n(x)$, $x \in X$. The integer $n(x)$ is the “number of balls of color x in the urn n .” The set of such urns on X is denoted by \mathcal{N} , formally:

$$\mathcal{N} = \mathbb{N}_+^X.$$

To each $n \in \mathcal{N}$ is associated the probability distribution \tilde{n} on X defined by

$$\tilde{n}(x) = \frac{n(x)}{\sum_{y \in X} n(y)}.$$

When we write that the alternative x is picked in the urn n , we mean that x is picked in X according to the probability \tilde{n} .

A *random urn sequence* is a sequence U_τ , $\tau \in \mathbb{N}$ of random variables on \mathcal{N} such that $U_{\tau+1}$ is defined conditionally on U_τ . Here are three examples:

1. Two-alternatives reinforcement. Given a realization $n_\tau \in \mathcal{N}$ of U_τ , an alternative x is picked in X according to the probability distribution $\tilde{n}_\tau^{[2]}$, and one ball of color x is added to the urn: $n_{\tau+1}(w) = n_\tau(w) + 1$ and for all $v \neq w$, $n_{\tau+1}(v) = n_\tau(v)$. This means that two alternatives, say a and b are picked independently in the urn n_τ , and are compared according to T . The result of the comparison is $x = \max\{a, b\}$, that is: $x = a$ if $a = b$ or if $a T b$ and $x = b$ if $b T a$. Alternative x is reinforced.

2. Three-alternatives reinforcement. Same thing as above, with the probability distribution $\widetilde{n}_\tau^{[3]}$. This means that three alternatives, say a , b and c are picked independently in X according to n_τ ; a , b and c are compared according to T in sequence and one ball of color $x = \max\{\max\{a, b\}, c\}$ is added to the urn. Remark that there are only two cases : ranked alternatives where we reinforce the top one or a cycle where we reinforce at random.
3. Fast reinforcement. Same thing as above, with the probability distribution $\widetilde{n}_\tau^{[\infty]}$, the stationary distribution for T when sampling is done according to \widetilde{n}_τ .

Remark that the first two examples can be concretely implemented easily, as described, but fast reinforcement cannot.

3.2 A motivational example

This section presents a simple non-rigorous argument to justify our focus on three-alternatives reinforcement. Consider the simplest possible non trivial tournament : a cycle of three alternatives A , B and C with ATB , BTC , CTA . In order to evaluate the long term behavior of two-alternatives and three-alternatives reinforcement, we use a deterministic continuous time motion corresponding to the limit of a large number of balls in the urn. We write $a(t)$, $b(t)$ and $c(t)$ the “number” of balls of each type and $\tilde{a}(t) = a(t)/t$, $\tilde{b}(t) = b(t)/t$ and $\tilde{c}(t) = c(t)/t$ the corresponding probabilities.

For two-alternatives reinforcement we get :

$$\begin{aligned} a' &= \tilde{a}^2 + 2\tilde{a}\tilde{c} \\ b' &= \tilde{b}^2 + 2\tilde{b}\tilde{a} \\ c' &= \tilde{c}^2 + 2\tilde{c}\tilde{b} \end{aligned} \tag{8}$$

and we remark that

$$\begin{aligned} \frac{d}{dt} (\ln \tilde{a} + \ln \tilde{b} + \ln \tilde{c}) &= \frac{d}{dt} (-3 \ln t + \ln a + \ln b + \ln c) \\ &= -\frac{3}{t} + \frac{\tilde{a}^2 + 2\tilde{a}\tilde{c}}{a} + \dots \\ &= -\frac{3}{t} + \frac{\tilde{a} + 2\tilde{c}}{t} + \frac{\tilde{b} + 2\tilde{a}}{t} + \frac{\tilde{c} + 2\tilde{b}}{t} \\ &= 0 \end{aligned}$$

so $(\tilde{a}, \tilde{b}, \tilde{c})$ cannot converge to the optimal probability independently of the state at finite time.

It is then natural to study three-alternatives reinforcement, for which :

$$\begin{aligned} a' &= \tilde{a}^3 + 3\tilde{a}^2\tilde{c} + 3\tilde{a}\tilde{c}^2 + \tilde{a}\tilde{b}\tilde{c} \\ b' &= \tilde{b}^3 + 3\tilde{b}^2\tilde{a} + 3\tilde{b}\tilde{a}^2 + \tilde{a}\tilde{b}\tilde{c} \\ c' &= \tilde{c}^3 + 3\tilde{c}^2\tilde{b} + 3\tilde{c}\tilde{b}^2 + \tilde{a}\tilde{b}\tilde{c} \end{aligned} \tag{9}$$

and for the same quantity

$$\begin{aligned} \frac{d}{dt} (\ln \tilde{a} + \ln \tilde{b} + \ln \tilde{c}) &= -\frac{3}{t} + \frac{\tilde{a}^2 + 3\tilde{a}\tilde{c} + 3\tilde{c}^2 + \tilde{b}\tilde{c}}{t} + \dots \\ &= \frac{4(\tilde{a}^2 + \tilde{b}^2 + \tilde{c}^2) + 4(\tilde{a}\tilde{c} + \tilde{b}\tilde{a} + \tilde{c}\tilde{b}) - 3}{t} \\ &= \frac{\tilde{a}^2 + \tilde{b}^2 + \tilde{c}^2 - 2(\tilde{a}\tilde{c} + \tilde{b}\tilde{a} + \tilde{c}\tilde{b})}{t} \end{aligned}$$

and simple calculus shows that this last term is positive expect for $\tilde{a} = \tilde{b} = \tilde{c} = 1/3$. Then $\ln \tilde{a} + \ln \tilde{b} + \ln \tilde{c}$ is an increasing negative function so it converges. It is not difficult to see, using the divergence of $\int 1/t dt$, that this implies that $\tilde{a}^2 + \tilde{b}^2 + \tilde{c}^2 - \tilde{a}\tilde{c} - \tilde{b}\tilde{a} - \tilde{c}\tilde{b}$ converges to 0 and then that $(\tilde{a}, \tilde{b}, \tilde{c})$ converges to $(1/3, 1/3, 1/3)$ (the details of the arguments will be given in the rigorous proof of the next section).

With this example we can see that two-alternatives reinforcement should not converge to the optimal probability even for a simple tournament and when we neglect the effect of probabilistic noise while three-alternatives reinforcement seems to work in that case. In the next section we will prove that three-alternatives reinforcement actually converges for any tournament. We will use the same idea of computing the variation of $\ln \tilde{a} + \ln \tilde{b} + \ln \tilde{c}$ with technical changes for the general tournament, the discrete time and the probabilistic evolution.

3.3 Three-alternatives reinforcement and martingale technique

We will now prove the result about three-alternatives reinforcement:

Theorem 4 *For any initial urn $n_0 \in \mathcal{N}$, the random urn sequence obtained by three-alternatives reinforcement is such that the realization n_τ , $\tau \in \mathbb{N}$ almost surely verifies:*

$$\lim_{\tau \rightarrow \infty} \widetilde{n}_\tau = p^*.$$

The same proof will actually also give the first part of result about two-alternatives reinforcement, which we thus state now:

Theorem 5 *For any initial urn $n_0 \in \mathcal{N}$, the random urn sequence obtained by two-alternatives reinforcement is such that the realization n_τ , $\tau \in \mathbb{N}$ almost surely verifies:*

$$\forall x \in X, p^*(x) = 0 \Rightarrow \lim_{\tau \rightarrow \infty} \widetilde{n}_\tau(x) = 0.$$

The proof relies mainly on the study of a well chosen function of the state of the urn. Let LD denotes a discrete logarithm: for integers $0 < a < b$,

$$LD[a, b] = - \sum_{i=a}^{b-1} \frac{1}{i}. \quad (10)$$

Recall that at time $\tau \in \mathbb{N}$, $n_\tau(w)$ denotes the number of w -balls in the urn. The total number of balls is increasing by 1 at each time, so $\sum_w n_\tau(w) = A + \tau$. The probability of drawing a w -ball is $\tilde{n}_\tau(w) = n_\tau(w)/(A + \tau)$. Consider the quantity

$$\mu_\tau = \sum_{w \in X} LD[n_\tau(w), A + \tau] \cdot p^*(w), \quad (11)$$

that is the expected value, according to the optimal probability p^* , of the discrete logarithm at time τ .

Proposition 6 *For both two-alternatives and three-alternatives reinforcement, the sequence μ_τ , $\tau \in \mathbb{N}$ is a negative sub-martingale. More precisely we have, for two-alternatives reinforcement:*

$$\mathbb{E}[\mu_{\tau+1} - \mu_\tau \mid n_\tau] = \frac{g(p^*, \tilde{n}_\tau)}{A + \tau}.$$

and for three-alternatives reinforcement:

$$\begin{aligned} & \mathbb{E}[\mu_{\tau+1} - \mu_\tau \mid n_\tau] \\ &= \frac{1}{A + \tau} \left(\frac{1}{2} g(p^*, \tilde{n}) + \frac{1}{2} \sum_{w \in X} g(w, \tilde{n})^2 p^*(w) + \sum_v \tilde{n}(v) g(p^*, v) (1 + g(v, \tilde{n})) \right). \end{aligned}$$

Proof. We will write p for \tilde{n}_τ and let i denote either 2 or 3. From τ to $\tau + 1$, one and only one ball is added. This ball has type w with probability

$p^{[i]}(w)$. Thus:

$$\begin{aligned}
& \mathbb{E}[\mu_{\tau+1} - \mu_{\tau} \mid n_{\tau}] \\
&= - \sum_{w \in X} p^{[i]}(w) \left(\sum_{v \neq w} \frac{1}{A + \tau} \cdot p^*(v) + \left[\frac{1}{A + \tau} - \frac{1}{n_{\tau}(w)} \right] \cdot p^*(w) \right) \\
&= \frac{-1}{A + \tau} + \sum_{w \in X} p^{[i]}(w) \frac{1}{n_{\tau}(w)} p^*(w) \\
&= \frac{-1}{A + \tau} + \sum_{w \in X} \frac{p^{[i]}(w)}{p(w)} \frac{p^*(w)}{A + \tau}
\end{aligned}$$

Where in the last line we used the definition $p(w) = n_{\tau}(w)/(A + \tau)$. Using the formula for $p^{[i]}$ of lemma 2, it comes, for two-alternatives:

$$\begin{aligned}
& \mathbb{E}[\mu_{\tau+1} - \mu_{\tau} \mid n_{\tau}] \\
&= \frac{-1}{A + \tau} + \sum_{w \in X} (1 + g(w, p)) \frac{p^*(w)}{A + \tau} \\
&= \frac{g(p^*, p)}{A + \tau}.
\end{aligned}$$

which is always non-negative. Furthermore, $g(p^*, p) = 0$ implies, by the first part of proposition 1, that $\text{Supp}(p) \subseteq \text{Supp}(p^*)$.

For three-alternatives, we have:

$$\begin{aligned}
& \mathbb{E}[\mu_{\tau+1} - \mu_{\tau} \mid n_{\tau}] \\
&= \frac{-1}{A + \tau} + \sum_{w \in X} \left(1 + \frac{3}{2}g(w, p) + \frac{1}{2}g(w, p)^2 + \sum_v p(v)g(w, v)g(v, p) \right) \frac{p^*(w)}{A + \tau}.
\end{aligned}$$

one can re-arrange:

$$\begin{aligned}
& (A + \tau) \mathbb{E}[\mu_{\tau+1} - \mu_{\tau} \mid n_{\tau}] \\
&= \frac{3}{2}g(p^*, p) + \frac{1}{2} \sum_{w \in X} g(w, p)^2 p^*(w) + \sum_v p(v)g(p^*, v)g(v, p) \\
&= \frac{1}{2}g(p^*, p) + \frac{1}{2} \sum_{w \in X} g(w, p)^2 p^*(w) + \sum_v p(v)g(p^*, v)(1 + g(v, p)).
\end{aligned}$$

And all the terms in this sum are non-negative. The sum can be 0 only if both $\text{Supp}(p) \subseteq \text{Supp}(p^*)$, and $g(w, p) = 0$ for all w in the support of p^* , which implies $p = p^*$ by the uniqueness in Proposition 1. **Q.E.D.**

We are now able to prove the two results of the beginning of the section.
Proof of Theorems 4 and 5. We consider for this proof either two or three alternatives reinforcement. We have:

$$\mathbb{E}[\mu_\tau] = \mathbb{E}[\mu_0] + \mathbb{E} \left[\sum_{t=1}^{\tau} \mu_t - \mu_{t-1} \right] \quad (12)$$

$$= \mathbb{E}[\mu_0] + \mathbb{E} \left[\sum_{t=1}^{\tau} \mathbb{E} [\mu_t - \mu_{t-1} | \mu_{t-1}] \right]. \quad (13)$$

By Proposition 6, μ_τ is a negative sub-martingale, so it converges almost surely to an integrable random variable μ_∞ (see Corollary VII.4.1 and VII.4.2 in [26]). Furthermore, the right hand side is an increasing negative sequence so it converges to a finite value. In the left hand side, the sum is an increasing sequence of positive random variables so by the monotonous convergence theorem (Theorem II.6.1 in [26]) we can take the limit inside the expectation. Hence $\mathbb{E} [\sum_{t=1}^{\infty} \mathbb{E} [\mu_t - \mu_{t-1} | \mu_{t-1}]] = \lim \mathbb{E} [\mu_\tau] - \mathbb{E} [\mu_0]$ is finite and so $\sum_{t=1}^{\infty} \mathbb{E} [\mu_t - \mu_{t-1} | \mu_{t-1}]$ is almost surely finite.

Let $f^{[2]}(p) = g(p^*, p)$ and $f^{[3]}(p) = g(p, p^*) + \sum_{w \in X} g(w, p)^2 p^*(x)$. The simplex $\Delta(X)$ is embedded in \mathbb{R}^X so we use the L_∞ distance. With this distance, $f^{[i]}$ is continuous and $d(\widetilde{n}_\tau, \widetilde{n}_{\tau+1}) \leq \frac{1}{A+\tau}$ almost surely. Denote by $B(p, \eta)$ the ball of center p and radius η .

Now consider a single realization of the urn process. Since X is a finite set, $\Delta(X)$ is compact, so let \widetilde{n}_∞ be an accumulation point for \widetilde{n}_τ . We will show that necessarily $f^{[i]}(\widetilde{n}_\infty) = 0$. Looking for a contradiction, suppose $f^{[i]}(\widetilde{n}_\infty) > 0$.

Since $f^{[i]}$ is continuous, let $\epsilon, \eta > 0$ be such that $\forall p \in B(\widetilde{n}_\infty, \eta)$, $f^{[i]}(p) > \epsilon$. Let ϕ be a sub-sequence such that $\forall \tau$, $\widetilde{n}_{\phi(\tau)} \in B(\widetilde{n}_\infty, \eta/2)$ and $\phi(\tau+1) > (1+\eta)\phi(\tau)$. Then:

$$\sum_{t=0}^{\infty} \mathbb{E}[\mu_{t+1} - \mu_t | n_t] \geq \sum_{t=0}^{\infty} \frac{1}{2(A+t)} f^{[i]}(\widetilde{n}_t) \quad (14)$$

$$\geq \sum_{\tau=0}^{\infty} \frac{1}{2(A+\phi(\tau))} \sum_{t=\phi(\tau)}^{\lfloor (1+\eta/2)\phi(\tau) \rfloor} f^{[i]}(\widetilde{n}_t) \quad (15)$$

$$\geq \sum_{\tau=0}^{\infty} \frac{1}{2(A+\phi(\tau))} \lfloor (1+\eta/2)\phi(\tau) \rfloor \epsilon \quad (16)$$

The right hand side of the last line is infinite. We already proved that the left hand side is almost surely finite. This contradiction proves that all

the accumulation points of \tilde{n}_t are zeros of $f^{[i]}$. It follows that $f^{[i]}(\tilde{n}_t) \rightarrow 0$ almost surely. We have seen in the proof of Proposition 6 that this fact implies exactly the theorems. **Q.E.D.**

3.4 Two-alternatives reinforcement and variance estimates

In this section, we study in detail the two-alternatives reinforcement. The main idea is to study the variance of μ_∞ conditionally on the state at a large time t .

Theorem 7 *For any initial urn $n_0 \in \mathcal{N}$, the random urn sequence obtained by two-alternatives reinforcement is such that:*

1. *almost surely, for all alternatives x such that $p^*(x) = 0$, $\tilde{n}_\tau(x) \rightarrow 0$ when $\tau \rightarrow \infty$;*
2. *with positive probability, the realized sequence $\tilde{n}_\tau, \tau \in \mathbb{N}$ has no limit as $\tau \rightarrow \infty$;*
3. *if T is such that $\forall x, p^*(x) > 0$ (in other words, $BP(T) = X$) then, with probability one, $\tilde{n}_\tau, \tau \in \mathbb{N}$ has no limit.*

To simplify notation, in this section we let the process start at $\tau \neq 0$ so that τ always denote the number of ball in the urn (i.e. $A = 0$). We will also only consider two-alternatives reinforcement in this section. Recall the piece of notation $\text{Supp}(p^*) = BP$. Also recall from the last section that μ_τ is a negative submartingale so it has an almost sure limit μ_∞ . Let $\phi = \sum_{x \in BP} p^*(x) \log p^*(x)$ be the value of μ_∞ when \tilde{n}_τ converges to p^* .

The first point is the following variance estimate :

Lemma 8 *Let $\tau > 0$ and let $\epsilon_\tau(x) = p^*(x)/\tilde{n}_\tau(x) - 1$. We have*

$$\mathbb{E}[(\mu_{\tau+1} - \mu_\tau)^2 | \mathcal{F}_\tau] = \frac{1}{\tau^2} \sum_{x \in BP} \tilde{n}_\tau^{[2]}(x) \epsilon_\tau(x)^2 \quad (17)$$

Proof. This is a straightforward computation :

$$\mathbb{E}[(\mu_{\tau+1} - \mu_\tau)^2 | \mathcal{F}_\tau] = \sum_x \tilde{n}_\tau^{[2]}(x) \left(- \sum_y p^*(y) \frac{1}{\tau} + p^*(x) \frac{1}{n(x)} \right)^2 \quad (18)$$

$$= \sum_x \tilde{n}_\tau^{[2]}(x) \frac{1}{\tau^2} \left(-1 + \frac{p^*(x)}{\tilde{n}_\tau(x)} \right)^2 \quad (19)$$

$$= \frac{1}{\tau^2} \sum_x \tilde{n}_\tau^{[2]}(x) \epsilon_\tau(x)^2 \quad (20)$$

Q.E.D.

The main point is the factor $\frac{1}{\tau^2}$ that make the series of those terms summable (once a small control on ϵ is provided). Thus the variance of μ_∞ conditioned on \mathcal{F}_τ will be of order ϵ^2/τ . Thus with high probability μ_∞ will be close to μ_τ so that if μ_τ is far enough from ϕ , then $\mu_\infty \neq \phi$.

We will first consider the case where $BP \neq X$. In this case we have $\mathbb{E}[\mu_{\tau+1} - \mu_\tau \mid \mathcal{F}_\tau] = \frac{g(p^*, \tilde{n}_\tau)}{\tau} \geq g_0 \cdot \tilde{n}(BP^c) > 0$ (where $g_0 = \inf_{y \in BP^c} g(p^*, y)$) so we need an estimate of $\tilde{n}(BP)$.

Lemma 9 *Suppose that there exists $\pi \in (0, 1)$ such that, at each time τ the probability of adding a ball in BP^c is inferior to $\pi \cdot \tilde{n}_\tau(BP^c)$. Then we have for all $\tau \geq \tau_0$,*

$$\begin{aligned} \mathbb{E}[\tilde{n}_\tau(BP^c) \mid \mathcal{F}_\tau] &\leq \prod_{t=\tau_0+1}^{\tau} \frac{t + \pi - 1}{t} \mathbb{E}[\tilde{n}_{\tau_0}] \\ &\leq \left(\frac{\tau_0}{\tau}\right)^{1-\pi} \tilde{n}_{\tau_0}. \end{aligned}$$

Proof. The first line comes from a straightforward induction

$$\begin{aligned} \mathbb{E}[\tilde{n}_\tau] &= \frac{1}{\tau} \mathbb{E}[n_{\tau-1}] + \frac{1}{\tau} \mathbb{E}[n_\tau - n_{\tau-1}] \\ &= \frac{\tau-1}{\tau} \mathbb{E}[\tilde{n}_{\tau-1}] + \frac{1}{\tau} \mathbb{E}[\mathbb{E}[n_\tau - n_{\tau-1} \mid \mathcal{F}_{\tau-1}]] \\ &\leq \frac{\tau-1}{\tau} \mathbb{E}[\tilde{n}_{\tau-1}] + \frac{1}{\tau} \mathbb{E}[\pi \tilde{n}_{\tau-1}] \\ &\leq \frac{\tau+\pi-1}{\tau} \mathbb{E}[\tilde{n}_{\tau-1}] \end{aligned}$$

and for the second line we have

$$\begin{aligned} \log \prod_{t=\tau_0+1}^{\tau} \frac{t + \pi - 1}{t} &\leq \sum_{\tau_0+1}^{\tau} \log\left(1 + \frac{\pi - 1}{t}\right) \\ &\leq \sum_{\tau_0+1}^{\tau} \frac{\pi - 1}{t} \\ &\leq (\pi - 1) \log\left(\frac{\tau}{\tau_0}\right) \end{aligned}$$

Q.E.D.

Now we are able to prove the third point of the theorem.

Proof of theorem 7, case $BP^c \neq \emptyset$. Remark that for $\delta > 0$ small enough and τ_0 big enough, the set

$$S = \{n \in \mathcal{N} \text{ s.t. } n(X) \geq \tau_0 \text{ and } |\phi - \sum_x p^*(x) LD(n(x), n(X))| \leq \delta\}$$

verifies :

- $\forall n \in S, \forall x \in BP, |p^*(x)/\tilde{n}(x) - 1| \leq 1$
- $\forall n \in S, \forall x \in BP^c, 1 + g(x, \tilde{n}) \leq \pi$

for some $\pi > 1 - \inf_{x \in BP^c} g(p^*, x)$.

Let us assume that $\tilde{n}_{\tau_0} \in S$, (which clearly happens with positive probability). Let T be the first time after τ_0 such that $\tilde{n}_T \notin S$. T is a stopping time so $\mu_{\tau \wedge T}$ is still a submartingale, let us call it μ'_τ . Furthermore, up to time T , we have

$$\begin{aligned} \mathbb{E}[\mu_{\tau+1} - \mu_\tau \mid \mathcal{F}_\tau] &= \frac{g(p^*, \tilde{n}_\tau)}{\tau} \\ &\leq \frac{\tilde{n}_\tau(BP)}{\tau}, \end{aligned}$$

and thus

$$\begin{aligned} \mathbb{E}[\mu'_{\tau+1} - \mu'_\tau \mid \mathcal{F}_{\tau_0}] &\leq \frac{1}{\tau} \mathbb{E}[\tilde{n}_\tau(BP) 1_{\{T \leq \tau\}} \mid \mathcal{F}_{\tau_0}] \\ &\leq \tau_0^{1-\pi} \tau^{pi-2} \tilde{n}_{\tau_0}(BP). \end{aligned}$$

Now let μ'_∞ denote the almost sure limit of $\mu_{\tau \wedge T} = \mu'_\tau$. Remark that if $|\phi - \mu_\infty|(\omega) < \delta$ then $T(\omega) = \infty$ and thus $\mu_\infty(\omega) = \mu'_\infty(\omega)$. It is therefore enough to show that, with positive probability $\phi - \delta < \mu'_\infty < \phi$.

Note that μ' is a bounded submartingale, so it also converges in L^1 and L^2 toward μ'_∞ . Thus

$$\begin{aligned} \mathbb{E}[\mu'_\infty - \mu'_{\tau_0} \mid \mathcal{F}_\tau] &= \sum_{t=\tau_0}^{\infty} \mathbb{E}[\mu'_{t+1} - \mu'_t \mid \mathcal{F}_{\tau_0}] \\ &\leq \sum_{t=\tau_0}^{\infty} \tau_0^{1-\pi} t^{pi-2} \tilde{n}_{\tau_0}(BP) \\ &\leq C \tilde{n}_{\tau_0}(BP), \end{aligned}$$

and

$$\begin{aligned}
\text{Var}(\mu'_\infty \mid \mathcal{F}_{\tau_0}) &\leq \sum_{t=\tau_0}^{\infty} \text{Var}(\mu'_{t+1} \mid \mathcal{F}_{\tau_0}) \\
&\leq \sum_{t=\tau_0}^{\infty} \mathbb{E}[(\mu'_{t+1} - \mu'_t)^2 \mid \mathcal{F}_{\tau_0}] \\
&\leq \sum_{t=\tau_0}^{\infty} \frac{1}{t^2} \sum_{x \in BP} \tilde{n}^{[1]}(x) \epsilon_t(x)^2 \\
&\leq C \frac{1}{\tau_0}.
\end{aligned}$$

Finally consider a τ_0 large enough so that $\frac{1}{\tau_0} \ll \delta$. It is clear that with a positive probability, $\tilde{n}_{\tau_0} \ll \delta$ and $|\mu_{\tau_0} - \phi|$ is close to $\delta/2$. Under this event we see that $|\phi - \mu'_\infty|$ is a random variable with expectation close to $\delta/2$ and variance small with respect to δ so $|\mu'_\infty - \phi|$ has a positive probability to be in $[\delta/4, 3\delta/4]$ which proves the theorem. **Q.E.D.**

Now we turn to the case where $X = BP$. The idea will be similar, with Lemma 8 being the core argument. The main simplification comes from the fact that in this case $g(p^*, p) = 0$ for all p so μ is a martingale and Lemma 9 will no longer be needed. However, in order to prove that μ_∞ is almost surely different from ϕ , we will need an almost sure lower bound on $|\phi - \mu_\tau|$ which will come from a careful analysis of the difference between discrete and real logarithms. Finally since the almost sure bound that we will get will be much worse than the one we were able to have with positive probability, we will need to be more careful in our use of Lemma 8.

First recall the following well known approximation result :

Proposition 10 *Let $k > 0$, there exists a constant γ (Euler's constant) such that*

$$\log(k+1) + \gamma - \frac{1}{2} \sum_{i=k+1}^{\infty} \frac{1}{i^2} \leq \sum_{i=1}^k \frac{1}{i} \leq \log(k+1) + \gamma - \frac{1}{2} \sum_{i=k+2}^{\infty} \frac{1}{i^2} \quad (21)$$

Furthermore when k tends to infinity,

$$\sum_{i=k+1}^{\infty} \frac{1}{i^2} \sim \frac{1}{k} \quad (22)$$

This proposition implies the following corollary:

Corollary 11 *Let T be a tournament on the set X such that $BP = T$. There exists $c > 0$ such that, for any urn $n \in \mathcal{N}$,*

$$\sum_x p^*(x) LD(n(x), \tau) \leq \phi - \frac{c}{\tau}. \quad (23)$$

Furthermore, writing $\epsilon(x) = p^*(x)/\tilde{n}(x) - 1$, if we restrict ourselves to large enough τ (with ϵ staying bounded) the constant c can be taken as close as we want to :

$$\frac{|X| - 1 + \sum_{x \in BP} \epsilon(x)}{2} \quad (24)$$

Proof. This is a straightforward computation using the definition of LD . **Q.E.D.**

We also need a control of ϵ in term of μ

Lemma 12 *Almost surely, for any time τ*

$$\sum_{x \in BP} [\tilde{n}_\tau(x) + p^*(x)/2] \epsilon_\tau(x)^2 \leq \phi - \mu_\tau \quad (25)$$

Proof. We have

$$\phi - \mu_\tau \geq \sum_{x \in BP} p^*(x) [\log p^* - \log p] \quad (26)$$

$$\geq \sum_{x \in BP} p^*(x) (\epsilon_\tau(x) + \epsilon_\tau(x)^2/2) \quad (27)$$

$$\geq \sum_{x \in BP} p^*(x) \epsilon_\tau(x)^2/2 + \sum_{x \in BP} \tilde{n}_\tau(x) (1 + \epsilon_\tau(x)) \epsilon_\tau(x) \quad (28)$$

$$\geq \sum_{x \in BP} (p^*(x)/2 + \tilde{n}_\tau(x)) \epsilon_\tau(x)^2 \quad (29)$$

because $\sum_{x \in BP} \tilde{n}_\tau(x) \epsilon(x) = \sum_{x \in BP} p^*(x) - \tilde{n}_\tau(x) = 0$. **Q.E.D.**

Together, the last lemmas have the following consequence

Lemma 13 *Let $\tau_0 > 0$ large enough such that $|\phi - \mu_{\tau_0}| \geq \frac{5}{\tau_0 - 1}$, then there exists $\pi > 0$ such that, with probability at least π ,*

$$\forall \tau > \tau_0, |\phi - \mu_\tau| \geq \frac{1}{\tau_0}.$$

Proof. Let $d = |\phi - \mu_{\tau_0}|$ and let $T = \inf\{\tau > \tau_0 | \phi - \mu_\tau \geq 2d \text{ or } \phi - \mu_\tau \leq d/5\}$. T is a stopping time so $\mu'_\tau = \mu_{\tau \wedge T}$ is still a sub-martingale ; by definition it is also bounded so it converges almost surely and in all L^p . Let μ'_∞ denote its limit. We will show that $\mathbb{P}(\mu'_\infty \in (d/5, 2d)) > \pi > 0$ for some π . Remark that, since μ makes vanishing steps, we are only interested in the behaviour of μ close to ϕ , so we can restrict ourself to small ϵ .

As long as $\tau < T$, by definition we have $\phi - \mu_\tau \leq 2d$ and thus by Lemma 12, $\sum_x (p^*(x)/2 + \tilde{n}_\tau(x))\epsilon_\tau(x)^2 \leq 2d$. For ϵ small enough, this implies $\sum_x \tilde{n}_\tau^{[1]} \epsilon(x)^2 \leq 2d$ and thus, by Lemma 8:

$$\forall \tau > \tau_0, \mathbb{E}[(\mu'_{t+1} - \mu'_t)^2 | \mathcal{F}] \leq \frac{2d}{\tau^2}. \quad (30)$$

Summing up to infinity (recall that μ' has constant expectation):

$$\text{Var}(\mu'_\infty - \mu'_{\tau_0} | \mathcal{F}_{\tau_0}) = \sum_{\tau=\tau_0}^{\infty} \mathbb{E}[(\mu'_{\tau+1} - \mu'_\tau)^2 | \mathcal{F}_{\tau_0}] \quad (31)$$

$$\leq \sum_{\tau=\tau_0}^{\infty} \frac{2d}{\tau^2} \quad (32)$$

$$\leq 2d \frac{1}{\tau_0 - 1} \quad (33)$$

$$\leq \frac{2}{5} d^2 \quad (34)$$

where we used the hypothesis $d \geq \frac{5}{\tau_0 - 1}$ in the last line.

Remark that, since $\mu_{T \wedge \tau}$ is bounded, $\mathbb{E}[\mu'_\infty] = d$. Moreover notice that a random variable with expectation d which never takes value in the interval $(d/2, 2d)$ has at least variance $d^2/2$. Since $\text{Var}(\mu'_\infty | \mathcal{F}_\tau) \leq \frac{2}{5} d^2$, we have $\mathbb{P}(\mu' \in (d/2, 2d)) \geq 1/10$ and on this event, $T = \infty$ so $|\phi - \mu|$ has never reached $d/2 \geq \frac{1}{\tau_0}$. **Q.E.D.**

Proof of Theorem 7, case $BP = \emptyset$. First note that if \tilde{n}_τ converges, it has to be toward a fixed point. It is easy to see that the fixed points of the dynamics are exactly the optimal strategies p_Y^* corresponding to all subtournaments $Y \subseteq X$. (This of course includes $p^* = p_X^*$ itself.) For $Y \subsetneq X$, $\sum_x p_X^*(x) \log p_Y^*(x) = -\infty$ so by the using the Markov inequality on μ we see that convergence to those fixed points is impossible. Therefore we only have to rule out convergence towards p_X^* .

We first consider the case $|X| \geq 12$. Then by Corollary 11, for any τ large enough, μ_τ almost surely verifies the hypothesis of Lemma 13. Fix any

suitable τ_0 , by Lemma 13, $\mathbb{P}(\tilde{n}_\tau \rightarrow p^*) \leq \mathbb{P}(\exists \tau \geq \tau_0 | \phi - \mu_\tau| \leq 1/\tau_0) \leq 1 - \pi$. On the event that $|\phi - \mu_\tau|$ does reach $1/\tau_0$ at time τ_1 , we can use Lemma 13 at time τ_1 to get $\mathbb{P}(\tilde{n}_\tau \rightarrow p^*) \leq (1 - \pi)^2$. By induction we get $\mathbb{P}(\tilde{n}_\tau \rightarrow p^*) = 0$ which proves the theorem.

For the case $3 \leq |X| \leq 11$ (a non trivial tournament has at least 3 elements), consider any τ_0 large enough. By Corollary 11 we have $|\phi - \mu_{\tau_0}| \geq 1/\tau_0$. Let $T = \inf\{\tau > \tau_0 | |\phi - \mu_\tau| \leq 1/2\tau_0 \text{ or } |\phi - \mu_\tau| \geq \frac{5}{\tau_0}\}$. The event $\{T = \infty \text{ or } |\phi - \mu_T| \geq \frac{5}{\tau_0}\}$ has probability at least $1/9$ and if $|\phi - \mu_T| \geq \frac{5}{\tau_0}$ we can apply Lemma 13 at time T so the conclusion of Lemma 13 is still true with $1/\tau_0$ replaced by $1/2\tau_0$ and we can use the same induction as before to prove the theorem. **Q.E.D.**

3.5 Conclusion

We have found the behavior of learning process designed to discover the “best” alternatives in a tournament. Learning is achieved through the following idea. An alternative which is considered as “good” at some date is reinforced for the future in the sense that one (slightly, and less and less) increases the probability for this alternative to be considered: reinforcement updates the sampling, or “prior” probability. The test according to which an alternative is considered as a good one at time t rests on comparing a few randomly chosen alternatives.

We found a very different behavior between the processes where reinforcement occurs after sampling two or three alternatives. With three alternatives, the process converges almost surely to a well-defined limit that has a nice interpretation in term of the tournament game: it is the optimal strategy for this game. One can therefore say that this form of learning is “successful”. With two alternatives, the picture is more complicated. The learning process “succeeds” in finding the Bipartisan set (a set which has been argued to be more important in term of social choice than the numerical values of the optimal probabilities [14]), but not the optimal probabilities themselves. We conjecture that the almost sure non-convergence happens for all tournaments and not only when $BP = X$.

Acknowledgment

Thanks to Bastien Mallein for useful discussions in the early stage of this study.

References

- [1] Brandt, Felix, Maria Chudnovsky, Ilhee Kim, Gaku Liu, Sergey Norin, Alex Scott, Paul Seymour, and Stephan Thomasse (2011) A counterexample to a conjecture of Schwartz. *Social Choice and Welfare* forthcoming.
- [2] Chebotarev, P. T., Shamis, E. 1998. Characterizations of scoring methods for preference aggregation. *Annals of Operation Research* **80**: 299—332.
- [3] Chebotarev, P. T., Shamis, E. 2006. Preference fusion when the number of alternatives exceeds two: indirect scoring procedures. arXiv:math/060217v3 [math.OC]
- [4] Daniels, H. E. 1969. Round-robin tournament scores. *Biometrika* **56**: 295—299.
- [5] David, H. 1963. *The Method of Paired Comparisons*, Griffin’s statistical monographs and courses. Charles Griffins, London.
- [6] Fishburn, P.C. 1977. “Condorcet social choice functions“ *SIAM Journal on Applied Mathematics* **33**: 469—489.
- [7] David C. Fisher and Richard B. Reeves, 1995. Optimal strategies for random tournament games. *Linear Algebra and its Applications*, **217**: 83—85.
- [8] Fisher, D., Ryan, J. 1992. Optimal strategies for a generalized ‘Scissors, Paper and Stone’ game. *American Mathematical Monthly* **99**: 935—942.
- [9] Fisher, D., Ryan, J. 1995a. Tournament games and positive tournaments. *Journal of Graph Theory* **19**: 217—236.
- [10] Fisher, D., Ryan, J. 1995b. Probabilities within optimal strategies for tournament games. *Discrete Applied Mathematics* **56**: 87—91.
- [11] Hofbauer, J., Schlag, K. 2000. Sophisticated imitation in cyclic games. *Journal of Evolutionary Economics* **10**: 523—543.
- [12] Laffond, G., Laslier, J.-F., Le Breton, M. 1993. The Bipartisan set of a tournament game. *Games and Economic Behavior* **5**: 182—201.

- [13] Laslier, J.-F. 1997. *Tournament Solutions and Majority Voting*, Berlin: Springer-Verlag.
- [14] Laslier, J.-F. 2000. Aggregation of preferences with a variable set of alternatives. *Social Choice and Welfare* **17**: 241—246.
- [15] Laslier, J.-F. 2000. Interpretation of electoral mixed strategies. *Social Choice and Welfare* **17**: 247—267.
- [16] Levchenkov, V. S. 1992. Social choice theory: a new insight. Discussion paper, Institute of Systems Analysis, Moscow.
- [17] McKelvey, R. 1979. General conditions for global intransitivities in a formal voting model. *Econometrica* **47**: 1085—1112.
- [18] Moon, J.W. 1968. *Topics on Tournaments*. Holt, Rinehart and Winston, New York.
- [19] Moulin, H. 1986. Choosing from a tournament. *Social Choice and Welfare* **3**: 271—291.
- [20] Myerson, R. B. 1993. Incentives to cultivate favored minorities under alternative electoral systems. *American Political Science Review* **87**: 856—869.
- [21] Myerson, R. B. 1995. Analysis of democratic institutions: structure, conduct and performance. *Journal of Economic Perspectives* **9**: 77—89.
- [22] Pemantle, R. 2007. A survey of random processes with reinforcement. *Probability Surveys* **4**: 1—79.
- [23] Posch, M. 1997. Cycling in a stochastic learning algorithm for normal form games. *Journal of Evolutionary Economics* **7**: 193—207.
- [24] Rivest, Ronald.L and Emily Shen (2010) “An Optimal Single-Winner Preferential Voting System Based on Game Theory” http://people.csail.mit.edu/rivest/gt/latest_conf.pdf
- [25] Rubinstein, Ariel 1996. Why are certain properties of binary relations relatively more common in natural languages? *Econometrica* **64**: 343—355.
- [26] Shiryaev, A.N. 1995. *Probability. Graduate Text in Mathematics*, Springer.

- [27] Slutzki, G. , Volij, O. 2006. Scoring of web pages and tournaments — axiomatizations. *Social Choice and Welfare* **26**: 75—92.
- [28] Ushakov, I. A. 1976. The problem of choosing the preferred element: An application to sport games. In *Management Science in Sports* (R. E. Machol, S. P. Ladany, and D.G. Morrison, eds.) 153—161. Amsterdam: North-Holland.